

Using High-Performance Networks to Enable Computational Aerosciences Applications

Marjory J. Johnson

The Research Institute for Advanced Computer Science is operated by Universities Space Research Association, The American City Building, Suite 311, Columbia, MD 21044, (301)730-2656

Work reported herein was supported by DARPA and the National Aeronautics and Space Administration under Cooperative Agreement NCC 2-387 to the Universities Space Research Association (USRA).

Using High-Performance Networks to Enable Computational Aerosciences Applications

Marjory J. Johnson
RIACS
NASA Ames Research Center
Moffett Field, California
U.S.A.

Abstract. One component of the U.S. Federal High Performance Computing and Communications Program (HPCCP) is the establishment of a gigabit network to provide a communications infrastructure for researchers across the nation. This gigabit network will provide new services and capabilities, in addition to increased bandwidth, to enable future applications. An understanding of these applications is necessary to guide the development of the gigabit network and other high-performance networks of the future. In this paper we focus on computational aerosciences applications run remotely using the Numerical Aerodynamic Simulation (NAS) facility located at NASA Ames Research Center. We characterize these applications in terms of network-related parameters and relate user experiences that reveal limitations imposed by the current wide-area networking infrastructure. Then we investigate how the development of a nationwide gigabit network would enable users of the NAS facility to work in new, more productive ways.

1. Introduction

The objective of the U.S. High Performance Computing and Communications Program (HPCCP) [8] is the development of the technology needed to address fundamental scientific problems, called *grand challenges*, that require orders of magnitude more computational power than is currently available. Eight federal agencies, including the National Aeronautics and Space Administration (NASA), are sharing responsibility for this program.* NASA's participation is focused on grand challenges in computational aerosciences, earth and space sciences, and remote exploration and experimentation.

A major objective of the HPCCP is the development of scalable, parallel computer architectures that are capable of sustained teraflops performance. However, in order to address the grand challenges, it is expected that multidisciplinary teams of scientists will

*The eight federal agencies are Defense Advanced Research Projects Agency (DARPA), Department of Energy (DOE), NASA, National Science Foundation (NSF), Department of Commerce (DOC)/National Institute for Standards and Technology (NIST), DOC/National Oceanic and Atmospheric Administration (NOAA), Environmental Protection Agency (EPA), and National Institutes of Health (NIH)/National Library of Medicine (NLM).

be required to collaborate with one another and with remote supercomputing resources. Accordingly, one of the technological components of the HPCCP is the development of a gigabit backbone network to interconnect the nation's researchers, educators, and computing resources. In addition to increased bandwidth, the gigabit network will provide new services and capabilities that will enable the collaborative way of conducting science that is envisioned for the future.

An understanding of application requirements is prerequisite to the development of future high-performance networks, such as the U.S. gigabit network. Scientific computing is considered to be one of the driving applications for such networks, because of the massive volumes of data generated by supercomputer simulations. In this paper we examine the role of networking to support computational aerosciences applications, in which computer simulation is used to predict aerodynamic characteristics of aircraft and aerospace vehicles. We focus, in particular, on computational fluid dynamics (CFD) applications.

Computational aerosciences is a grand-challenge area of particular importance to NASA Ames Research Center, since it is coordinating NASA activities in this area. NASA Ames has been a center for aerodynamics research for many years; some of the world's largest wind tunnels are located at this site. In recent years the value of computational aerosciences (including CFD) in complementing wind-tunnel experimentation has become widely accepted. An overview of CFD and its role in designing aircraft and aerospace vehicles via supercomputer simulation is presented by Vaziri [9] and Bailey [2]. NASA Ames' Numerical Aerodynamic Simulation (NAS) facility, which became fully operational in 1987, is a supercomputer center dedicated primarily to computational fluid dynamics. In effect, aircraft configurations are tested by "flying" them in the NAS supercomputer system, thus reducing the amount of wind-tunnel testing and experimental flight testing that is required.

The massive amount of data generated by CFD simulations places enormous demands on both the local-area and wide-area networking infrastructure. Studies in the literature that address data-communication requirements to support CFD include a 1988 study by Levin et al [6], which projects requirements for local-area network bandwidth at the NAS facility in the early 1990's, based on expected increases in computer speed; a 1990 paper by Lekashman [5], which describes typical CFD problems and their associated data-transfer requirements; and a 1991 paper by Vaziri [9] which presents data requirements for scientific visualization.

In this paper we focus on wide-area communications requirements to support remote users of the NAS facility. We examine the impact of the current wide-area networking infrastructure on their work, discuss potential future applications that would be enabled by the existence of higher performance wide-area networks (such as the gigabit network being developed in the U.S. as part of the HPCCP), and discuss network services that would be required to support these applications.

2. Using the NAS Facility

The NAS facility was established to support research and development for commercial and military aircraft. The facility contains two supercomputers, a Cray-2 and a Cray Y-MP, which are accessed by users across the nation.

In this section we relate experiences of computational fluid dynamicists who access the NAS facility via wide-area networks. We spoke with users at NASA Langley Research Center (LaRC), Hampton, Virginia; NASA Lewis Research Center (LeRC), Cleveland, Ohio; McDonnell Douglas, St. Louis, Missouri; and Rockwell International, Los Angeles, California. Besides being some of the heaviest users of the NAS facility, these users are also representative of the NAS user community in several ways. First, these locations span the U.S. In addition some of these users are connected to the NAS facility via multiple T1 links (where a single T1 link provides 1.544 megabit/second access), while others have only 56 kilobit/second access. Finally, some of the users we spoke with are researchers, while others are actively involved in aircraft development.

Users were asked to explain how they use resources at the NAS facility, to characterize their applications, to relate networking problems they have encountered in using the NAS facility, to describe the kinds of applications they envision for the future, and to describe how they collaborate with others in their NAS-related work. The first three issues are addressed below; the last two issues are addressed in Section 3. To establish a background for these discussions, we begin by identifying NAS users and describing the network infrastructure which supports them.

2.1. Profile of Remote User

The NAS facility is a national resource. Since it became fully operational in 1987, the number of users has increased steadily. At present there are approximately 1500 users at over 150 sites, including NASA centers, Department of Defense laboratories, aerospace industries, and universities. During normal working hours there are typically between 150 and 200 simultaneous users of the NAS supercomputers.

Several of the users that we contacted said that they use the NAS facility daily. Though supercomputer resources often are also available at their home site, they use the NAS facility because their local facilities provide insufficient memory to support their applications. We received conflicting answers when we inquired about past trends and future projections for using the NAS facility. Some, such as users at McDonnell Douglas, said their use of the NAS facility has increased over the years and is likely to continue to increase. On the other hand, users at NASA LeRC said that their use of the NAS facility (which initially increased steadily) has leveled off and is likely to decrease next year because they have recently acquired an on-site supercomputer that has adequate memory to support some of their applications.

Remote users typically use the same desk-top tools as the local users. In particular they use the same kind of workstation and the same software visualization packages, which in fact were developed by personnel at NASA Ames. We discuss the role of visualization and workstations in more detail in section 2.3.

2.2. Network Infrastructure to Support NAS Users

Recently the local area network at the NAS facility was upgraded from a combination of HYPERchannel and Ethernet to a combination of UltraNet, FDDI, and Ethernet. The UltraNet provides essentially gigabit local-area network capabilities for on-site users.

Remote users access the NAS facility through various national networks. Users at other NASA sites and at aerospace industries, i.e., the heaviest off-site users, access the facility via a NAS-provided network called AEROnet. Users at government laboratories typically access the facility via the Department-of-Defense-sponsored MILnet, while users at universities typically access it via the Internet. AEROnet is illustrated in Figure 1. The AEROnet backbone network consists of dedicated T1 links, with multiple T1 links to designated NASA sites. Leaf sites are connected to backbone sites via dedicated 56 kilobit/second links.

The original wide-area network servicing users at other NASA sites and at aerospace industries was called NASnet; the installation of AEROnet was completed at the end of 1991. One improvement of AEROnet over NASnet was the upgrade from a single T1 link to multiple T1 links to designated sites. For a detailed description of AEROnet, see [1].

2.3. Role of Visualization and Graphics Workstations

Massive volumes of data are generated by computational aerosciences simulations; it is virtually impossible to comprehend such large amounts of raw data without the aid of computer visualization, i.e., using graphics and imaging techniques to visualize fluid flow. Hence, visualization is a basic tool for the computational aer scientist. Theoretically, visualization can either be done interactively with the supercomputer or via a workstation at the user site. However, both insufficient network bandwidth and scarce supercomputer cycles (in demand for other tasks that require more compute power) make the latter method more practical at the present [7].

Both local and remote users of the NAS facility use personal high-performance graphics workstations to interface to the NAS supercomputers. The raw data produced by a supercomputer simulation, called the solution file, is typically transferred to the user site for post-processing, analysis, and visualization of the results on the scientist's workstation. NAS personnel have developed software packages that enable real-time visualization at the user's workstation. For example, PLOT3D, widely used by everyone we contacted, is an interactive graphics program developed by NAS personnel that displays two-dimensional and three-dimensional CFD data sets. The fact that PLOT3D, which must be resident at the user's workstation, is designed for a particular vendor's equipment, means that remote users' workstations are typically the same (or, perhaps earlier versions of the same machine) as those provided for on-site users.

AEROnet User Sites

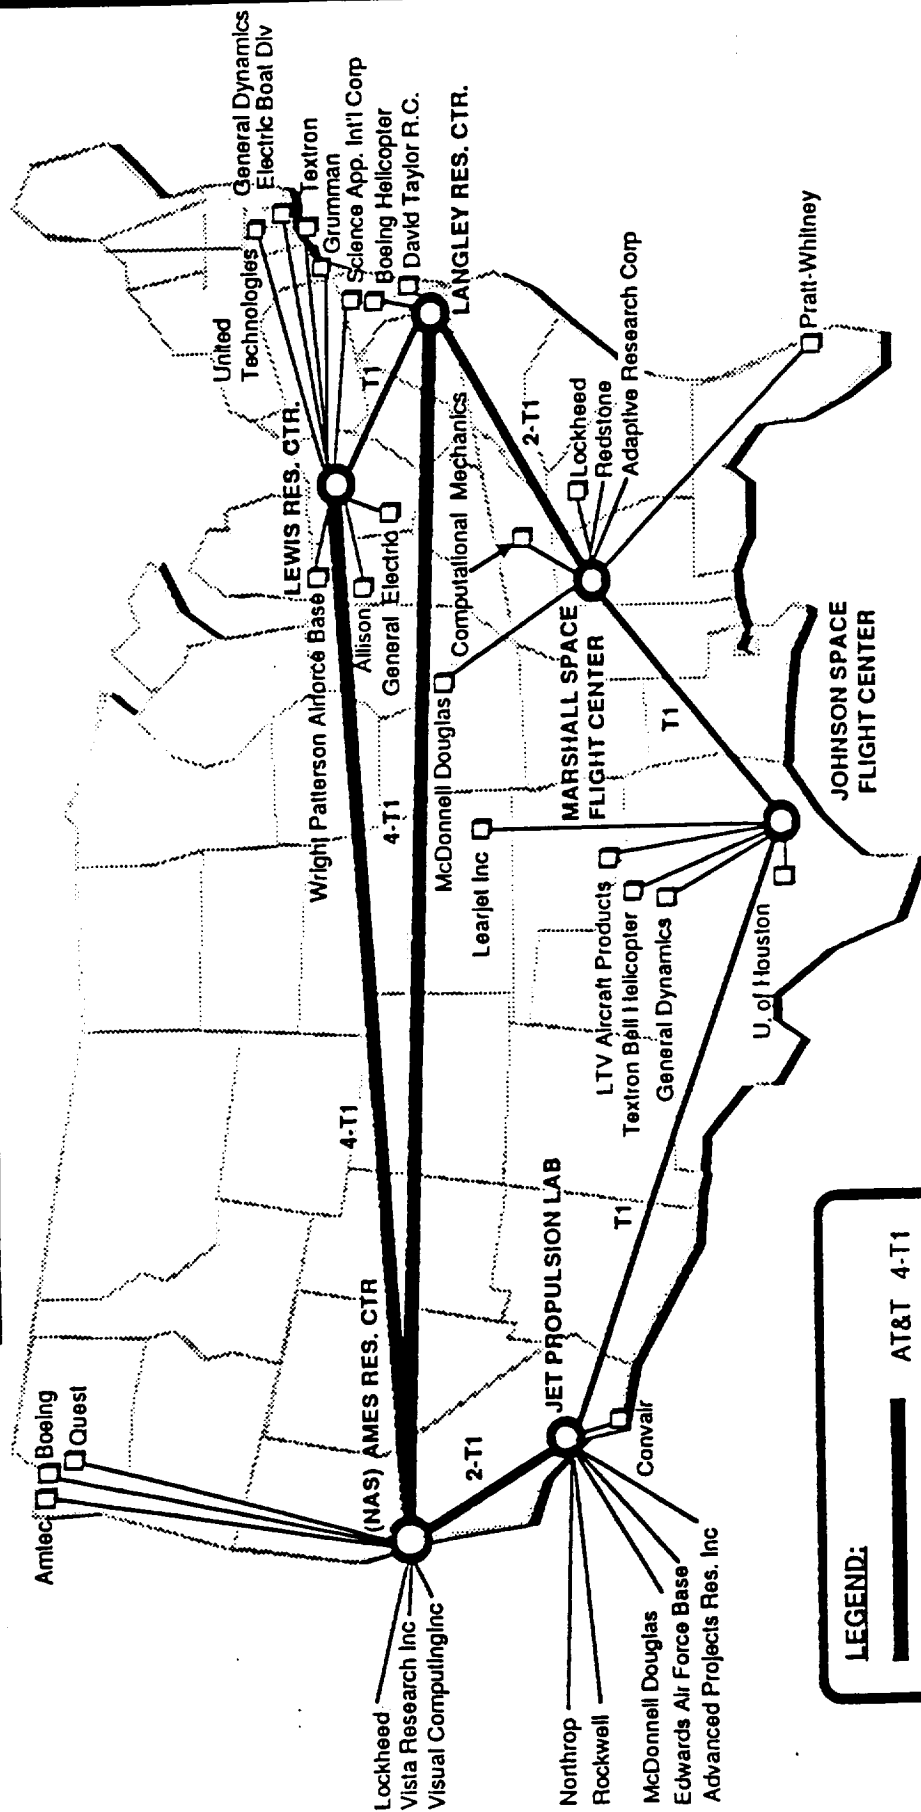


Figure 1.

2.4. Description of Current Applications

As indicated earlier, the NAS facility is dedicated primarily to CFD. In this section we give a brief introduction to CFD, discuss the general procedure for using the NAS facility remotely, and characterize user applications in terms of parameters that influence network performance from the user's perspective.

Fluid flow is described theoretically by the Navier-Stokes equations. In complex situations it is impossible to obtain exact solutions to these equations; simulation is used instead to find iterative solutions to approximations of the Navier-Stokes equations. The first step is to generate a grid to cover the structure and flow field that is being modeled. The resulting grid points are used in computations during the simulation. Hence, the size of the job, both in terms of computer resources and network resources that are required, depends both on the size of the structure being modeled and the granularity of the grid.

Simulations can either model static behavior, i.e., behavior at a single instance of time, or they can model dynamic behavior, i.e., behavior at successive time instances. Dynamic modeling is becoming increasingly significant, since it produces a more realistic representation of aerodynamic behavior. Time-accurate solution files, produced when modeling dynamic behavior, can be quite large, since they include a solution file for each time instance that is modeled. One user said that time-accurate analyses that he is currently conducting can include anywhere from 1000 to 10000 time instances.

Computational fluid dynamicists typically run their simulations in batch mode. However, they work interactively on computers at the NAS facility to edit source code, monitor their jobs, and do some post-processing of solution files before the data is transferred back to the user site. One user indicated that most of the researchers at his organization typically keep a workstation window open on one of the NAS computers for interactive access.

Relevant files associated with the batch jobs are the source code, the grid file, and the solution file. The source code and the grid file are prepared at the user site and transferred once to the NAS facility. Later modifications or editing are done interactively. The solution file is transferred back to the user site after the simulation is completed. The solution file, in particular, is often very large. Transfer of large CFD data files can tax the bandwidth capabilities of the current networking infrastructure.

Simulation runs typically take several hours to complete. Various techniques are used to divide the runs into segments two to three hours long. Users interactively monitor their jobs to check on status and/or to ensure that the program is producing reasonable results (e.g., results that are converging). Some users indicated that they always keep a window open on one of the machines at the NAS facility to watch the job queue to determine when their job is running. Other users described looking at data produced by the simulation every two or three hours. If this data indicates that there is a problem, then they interactively modify the source code or grid file and resubmit the job.

At the end of a run, some post-processing may be done before the solution file is transferred back to the user site. For example, the solution file may be split into smaller segments before transfer, or only certain portions of the solution file may be selected for

transfer. Such post-processing is done to reduce the size of the file that must be transferred over the networks.

Specific applications that the users described to us varied from modeling a single aircraft component to modeling an entire aircraft. As a result, there is a wide range of file sizes and run times associated with the various applications. The largest static-behavior model that was described is a full configuration of the F-18 aircraft, including part of an inlet duct. There are 3.7 million grid points for the application, the grid file is 160 megabytes, run time is approximately 30 to 50 hours, and the solution file is 170 megabytes.

In contrast to the large file sizes associated with the F-18 application, one user specified an average solution-file size of 2 to 3 megabytes. Lekashman [5] discusses steady-state problems whose solution files range in size from 4 kilobits to 32 megabytes, and a time-accurate problem with 1000 time instances. In addition he specifies 1 million grid points as being a typical number. A comparison of Lekashman's data, which is about a year old, to data collected for this study indicates that users continue to work on increasingly bigger problems that generate more and more data and place increasing demands on the network infrastructure.

Typical run times stated for applications other than the F-18 aircraft include approximately 10 hours for one problem, and 3 hours per time instance for another problem. All of the users said that they divide their problem into segments which require 2 to 3 hours to run. Simulation run time, of course, places no demands on the wide-area networks servicing the users. However, awareness of the length of the runs adds valuable perspective when trying to determine a reasonable time frame within which a user's data should be delivered to his home site after a run is completed. That is, users who are accustomed to waiting for their runs are likely to be willing to wait for their results. It is interesting that some users indicated they weren't sure exactly how much time their runs require. Their normal working procedure is to find an idle terminal, submit their job, go off to do something else (maybe even leaving overnight), and periodically come back to check on the run.

Network transfer time varies both with the size of the file being transferred and with the speed of the network connection. At NASA LaRC, which is connected to the NAS facility by 4 dedicated T1 lines (providing 6.176 megabit/second access), a user said transfer time was typically within 1 minute. At NASA LeRC, which enjoys the same access to the NAS facility as NASA LaRC, a user said transfer time was typically within 10 minutes. At McDonnell Douglas, which has only 56 kilobit/second access, they said transfer of solution files for the F-18 application can take from 4 to 8 hours; they said they really don't know for sure, because they don't wait around and watch.

2.5. Coping with Network Deficiencies

The major frustration expressed by users of the NAS facility is turn-around time for their simulations, because their jobs must wait in the queue for such a long time before they are run. Most of the users spoke of increased turn-around time over the years, because of increased usage of the NAS facility.

Limited bandwidth is the major cause of network-related user frustrations. In fact, limited bandwidth is one of the reasons why visualization of simulation results is done at the user's site, rather than with the supercomputer.

Users uniformly experienced frustration when transferring large data sets across the network. A common observation was that if you need to transfer a large data set, the link is sure to go down in the middle of the transfer. This complaint came from users with multiple T1 access, as well as users with 56 kilobit/second access. However, the latter said they were sure they wouldn't experience so much frustration if they had T1 access. A user at NASA LeRC, who said that typical solution-file transfer time is approximately 10 minutes, said that it can take hours if there is a problem with the link. Link reliability is clearly a related issue. An extreme example of this data-transfer/link-reliability problem involves McDonnell Douglas. When we first contacted them, they were expecting a call from NAS user services; they were going to request that their data be mailed to them on tape, because they were so frustrated in trying to transfer it over the network. Apparently mailing a tape is not uncommon for users with 56 kilobit/second access. One such user indicated they do this about twice a year, when transferring large files to the NAS facility.

Users have generally learned to cope with the problem of transferring large data sets by dividing them into smaller segments before attempting network transfer. This was done by all the users we contacted.

Network response time is a sporadic problem for some users. For example, one user, who has 56 kilobit/second access to the NAS facility, related frustrating delays during interactive editing sessions using *vi*. He said that problems of this nature occur maybe every three weeks, and last for several days in a row before the problem is corrected. However, in the context of the current mode of operation, network response time is not nearly as critical a parameter for the user as network bandwidth.

All the users we contacted said they were satisfied with their current mode of interaction with the NAS facility. Users have successfully devised techniques for coping with current networking limitations and are satisfied with the status quo. In fact, it seems to take considerable energy to introduce new modes of operation. As one user remarked, users shy away from having to learn a new way of doing things. He indicated that a guinea pig is needed, who is willing to take the time to learn the new techniques and then convince the other users that it really isn't so hard, after all. Despite this attitude, all the users we contacted responded enthusiastically when we described the more collaborative ways of working that are presented in the next section.

3. Future Applications

The obvious immediate benefit of higher-speed (e.g., gigabit) access to the NAS facility is faster transfer of large files. However, probably even more significant is the new way of doing science that gigabit networks will enable. In this section we address the increased data-transfer requirements of future applications and discuss various types of collaborative work that will be enabled by the establishment of a nationwide gigabit-network infrastructure.

3.1. Increased Data-Transfer Requirements

As we have seen, solution files are already large enough to place enormous demands on network bandwidth. One user, who currently has 56-kilobit access to the NAS facility, said it is impossible for his organization to do 3-D dynamic analyses at present, because of insufficient network bandwidth to return the resultant solution files. When asked what kinds of applications users envisioned for the future, they talked of transient-oscillation and turbulence problems, multi-disciplinary studies, and analyses using more grid points. All these types of studies will place increasing demands on processing power, memory capacity, and data-transfer capabilities.

Technological advances that are made in the development of more powerful computers will enable more sophisticated and complex CFD models to be developed. Finer grid granularity will be possible, models of complete aerospace vehicles will become possible, and time-accurate simulations will be more prevalent. All these factors will contribute to the increasing size of solution files, intensifying the need for higher-speed access to the NAS facility.

One user at an aerospace industry, involved with modeling the full configuration of the F-18 aircraft, offered the following explanation for recent growth in solution-file size for static-behavior analyses. He indicated that eighteen months ago, solution files were about a quarter of their current size. He attributed the increase in size to advances in grid-generation techniques. He reasoned that if a grid can be generated in a month, someone will pay for the work to be done; if grid generation takes 6 months, they won't. He said he doesn't expect a further increase in solution-file size for about 6 months.

3.2. Interactive Visualization

Currently visualization of simulation results is done at the user site after the simulation run is complete. In the future an interactive mode of visualization may be possible, i.e., while a simulation is in progress, the user may be able to change parameters and visualize the results on his workstation screen immediately. This is called *steering* a simulation.

While the users expressed satisfaction with their current mode of operation, they all thought interactive visualization would be beneficial. This type of real-time interactive activity would provide an attractive alternative to examining numerical solution-file data every two to three hours during a simulation.

The network infrastructure in place today is too slow to support interactive visualization. Transmission of images is extremely bandwidth intensive. According to Vaziri [9] an image data set is 1024 X 1280 pixels/frame with 4 bytes/pixel, to give a total of 5.24 megabytes/frame. Animation at a rate of 12.5 frames per second, which Vaziri calls a motion threshold, requires 65.5 megabytes/second network bandwidth. Flicker-free animation would require 60 frames/second, or 314 megabytes/second. Even a network transfer rate of 1 gigabit per second would be inadequate to handle the latter load without using data compression to reduce the volume of data.

3.3. Scientific Collaboration

A working hypothesis within the HPCCP is that teams of scientists will be required to address complex scientific problems of the future. Such teams of scientists are likely to be located remotely from one another and remotely from the supercomputers they need to use. Interactive, computer-supported collaboration will be required to enable this type of activity.

The kind of interactive collaboration envisioned as part of the HPCCP would represent a new way of working for the users we contacted. Some said they don't really collaborate at all in their work; others said they collaborate only with scientists within their organization; still others said they collaborate with people outside their organization as well. One user said that sometimes his organization sends a person to NASA Ames to work with on-site personnel. Collaborative tasks in which the users commonly participate include designing algorithms, coding, analyzing results, and writing joint papers. The mechanisms used to support these collaborative tasks are telephone calls and face-to-face meetings (including project-planning and review meetings, as well as working meetings), along with sharing of computer data files.

When asked about the desirability of collaborative, interactive work while using the NAS facility, they were all extremely enthused. One user said he thought it was valuable for two scientists to be able to see the same picture on their workstation screens during a technical discussion, so that they could easily refer to points of interest. However, when asked if he would find it useful to include voice capabilities in computer-supported collaborative activities, he replied that it would be disruptive in his working environment, because his desk isn't sufficiently isolated from other workers.

We described a scenario to the users that involved scientists in different locations simultaneously visualizing a simulation while it is running, with perhaps one scientist designated to steer the simulation. All the users were genuinely excited about being able to work in this manner. The number of people involved in a collaboration of this type was envisioned to be small, according to all the users questioned. One user suggested that future collaborations would probably involve two sites with two or three people at each site. Another user suggested three people, perhaps including one at a NASA center and two from industry or one from industry and representatives from two different NASA sites; still another suggested that a collaboration would involve two people most of the time and would probably never involve more than three.

A few suggestions were made for other uses of computer-supported collaboration. One suggestion, popular with managers, is to use this technique for monthly reviews. Another, more novel, suggestion is collaborative trouble-shooting involving a scientist and user-support personnel.

3.4. Multidisciplinary Modeling and Distributed Computing

Several disciplines play a role in aerospace vehicle design and analysis, including fluid dynamics, chemistry, structural dynamics, propulsion, solid mechanics, control theory, acoustics, thermodynamics, and combustion. According to Voigt [10] "the

traditional approach to the design of aerospace vehicles is to treat each discipline separately and in turn.” Both Bailey [2] and Voigt [10] indicate that as future aerospace vehicle designs are required to satisfy finer and finer tolerances, interactions between the disciplines listed above must be more tightly coupled.

Comments from the users we contacted confirmed that multidisciplinary modeling is assuming increasing importance. One user described a current application which consists of a tightly coupled family of codes, with information flowing from one section of the code into another. The run is divided into segments; results obtained from one segment are used as input for the next. Another user spoke of using object-oriented programming to link structures code and CFD code together.

From the above examples, it is clear that current multidisciplinary modeling is generally not interactive in nature. When asked about the desirability of interactive multidisciplinary modeling, with models developed for different disciplines running simultaneously and passing data back and forth, one user said this would be handy now, but would really be useful in their future study of transient-oscillation issues.

Interactive multidisciplinary modeling is likely to involve distributed computing. Though several users said that they currently use computers at different locations as they work on a problem, these computers are not working interactively. For example, one user at LaRC talked about using results from a run at NAS as input for a subsequent run at LaRC. As another example, a user described using computers simultaneously at both NAS and LaRC, but the work is split between the two sites in such a way that the pieces are completely stand-alone. Future multidisciplinary modeling will involve multiple computers working on a problem simultaneously, sharing data during the process. One user envisions a multidisciplinary study using different machines, perhaps at different geographical locations, dedicated to different disciplinary tasks, e.g., fluids and structures. At the same time, another machine might be used for readapting grids and perhaps still another machine might govern the overall activity via an expert system.

4. Networking Requirements to Support Future Applications

Services that have been envisioned for the HPCCP gigabit network, as presented in [8], include supporting “access to ... large scale distributed computing resources” and supporting “computationally intensive applications that require real time visualization of modeling and simulation results, rapid interrogation and retrieval of scientific data from specialized data bases, remote control of experiments and simulations, and teleconferencing.” These services are clearly required to enable the computational aerosciences applications presented in the previous section. In this section we investigate specific networking capabilities that are required, using information presented in Sections 2 and 3 to justify the requirements.

4.1. Bandwidth Requirements

Bandwidth requirements for the future will be driven both by image transfer and by transfer of large data sets. As we have seen from Section 2.4 solution-file size for some current applications already exceeds one gigabit, while from Section 3.1 it is clear that solution files will continue to increase in size in the future. Though users currently cope with bandwidth limitations by dividing solution files into segments before data transfer, it is preferable to maintain solution files (including time-accurate solution files) as single files. Also, with respect to image-transfer requirements, we have seen from Section 3.2 that interactive visualization can easily require multi-megabyte or gigabit bandwidth to support a single user. To compound the problem, when scientists are collaborating, several users (probably at different geographical locations) will need to view the same images simultaneously. Even though the use of data-compression techniques would reduce this requirement significantly, the bandwidth requirements are still overwhelming.

Since higher-speed links are prohibitively expensive to install as part of AEROnet at this time, NAS personnel have modified the *ftp* file transfer program to utilize available bandwidth more effectively [4,5]. The original Transmission Control Protocol (TCP) uses a sliding-window flow-control scheme; transmission is possible only while the window is open. With links that have a high delay-times-bandwidth product, with the capacity for a lot of data to be in transit at any given time, the window can easily fill up before any acknowledgements are received. Such a situation would cause data transmission to be interrupted and network bandwidth to be wasted. The modified protocol, called *mftp*, alleviates this problem by allowing the user to specify that multiple data streams be used in subsequent file transfers. The effective window size then becomes the actual window size times the number of streams.

The optimal number of data streams is clearly dependent on characteristics of the link being used. The recommendation in [4] is that 6 or fewer channels should be sufficient to keep data flowing in a dedicated coast-to-coast T1 link (e.g., NASA Ames to LaRC or NASA Ames to LeRC). However, processing delays introduced by gateways limit the improvement in rate of file transfer experienced by users who access the NAS facility via the Internet.

4.2. Response-time Requirements

Network response time is not currently a major concern for users. They have adjusted their mode of operation to the expected performance of the network; unless there is a problem with the link, they have no complaints. However, interactive visualization, collaborative activities, and distributed computing will all place stringent demands on network responsiveness.

For example, during interactive visualization, images would need to traverse the network in near-real time. A user would need to be assured of receiving network support for sustained high-rate data transfer during an interactive-visualization session. Because simulation runs are typically so long, it wouldn't make sense to reserve bandwidth to support interactive visualization during the entire run. However, it would be desirable to allow the user to schedule either an interactive-visualization session or a collaborative-

work session for a designated period of time. Then the network would need to be able to ensure the availability of adequate resources to support these activities as scheduled.

4.3. Coordination Requirements

In addition to low-response-time requirements, collaborative work requires synchronization between the collaborators. The network must support multicast to a small number of locations (probably two or three, according to Section 3.3), along with the ability to reserve the bandwidth resources to service all the users simultaneously. The amount of data that is likely to be involved is enormous, since users will want to share images, as well as textual data.

All collaborators in a collaborative session need to receive the same kind of service from the network infrastructure. A comparison of experiences from users with multiple T1 access and those with 56 kilobit/second access to the NAS facility indicates that users who receive noticeably different qualities of service from the network would not be able to participate effectively as part of the same collaborative team.

Coordination is also essential to support multidisciplinary modeling and distributed computing. Data transfers between the various machines in a distributed-computing system need to be coordinated, so that program segments are supplied with the proper data when they need it.

4.4. Requirements for Supporting Diverse Traffic Types

The future computational aerosciences applications presented in Section 3 involve traffic of many different types, each requiring different kinds of support from the network. For example, interactive visualization will require extremely high bandwidth on a steady basis during the entire visualization session, while transfer of a single image will generate a single high-rate burst of data on the network. Low and predictable delivery time is required to support both interactive visualization and collaborative work sessions, while immediate delivery of a single image is not so critical. Also, we note that although users are generally eager to incorporate interactive visualization, scientific collaboration, and multidisciplinary modeling into their working environment, the kinds of traffic associated with current activities must still be supported by the network. For example, the transfer of large solution files will still be important. Also, users will continue to expect support for electronic mail and for the interactive editing of files. Network protocols of the future must be able to provide all of these types of traffic with the services they require.

Though current networks already handle varying traffic types, the differences between types (and the services they require from the network) will be magnified as data sets become larger and response-time requirements become more stringent. According to McCormick et al [7], networking must undergo a fundamental change in order to support interactive visualization and modeling adequately. The claim is made that the network must be a visualization-based network, rather than being optimized for efficient transmission of textual data, and that new protocols must be developed accordingly.

Some actual measurements of traffic at the NAS facility, made using a program written by Greg Chesson [3] to monitor computer traffic, illustrate the current mix of traffic handled by the NAS networking infrastructure. Table 1 presents statistics that were collected over a 20-minute time period on a user-oriented local-area network at the NAS facility. For each packet length (or range of lengths if the number of packets of each individual length in this range accounts for less than 1% of the total bandwidth), Table 1 lists the percentage of the total number of packets that are of that length and the percentage of the total bandwidth that is used by those packets. Note the large number of minimum-length 60-byte TCP/IP packets (which constitute a small percentage of the overall bandwidth) and the small number of very large packets (over 8000 bytes), which constitute a significant percentage of the overall bandwidth. Based on user information, it is likely that these large packets represent solution-file transfers from one machine to another, or to mass storage. For example, one user described moving a solution file to a particular machine at the NAS facility for post-processing before transferring the data back to his home site.

Although Chesson warns that many of the packets that are less than 200 bytes long are either acknowledgements or control packets, it is clear that interactive traffic is also a significant percentage of overall bandwidth. In fact Chesson indicates that the statistics collected on the NAS network follow the same general pattern of statistics collected on general-purpose networks.

Table 1. Traffic Measurements

Packet Length (Bytes)	% of Total Packets	% of Total Bandwidth
60	41.4	6.6
61-97	7.0	1.4
98	4.4	1.2
99-125	4.8	1.4
126	5.9	2.0
127-137	1.4	.5
138	4.2	1.5
139-565	12.2	6.1
566	5.7	8.5
567-1077	2.6	5.8
1078	5.2	14.8
1079-1513	.5	1.8
1514	2.5	10.1
1515-8299	.7	7.6
8300	1.1	24.8
8312	.3	5.6

5. Conclusion

High-performance networks, in conjunction with the development of more powerful computers and improvements in visualization techniques, will enable computational aerosciences applications of the future. In this paper we examined the networking infrastructure which provides remote access to the NAS facility located at NASA Ames Research Center. We presented recent experiences of some of the current users and we discussed possible future CFD applications. Finally, we derived networking requirements to support these future applications.

AEROnet, the wide-area network which provides access to the NAS facility for its heaviest off-site users, supports a homogeneous environment. Since it is a special-purpose network, dedicated to the support of the NAS facility, application requirements can be specified in detail. Also, users all have the same local working environment; they conduct their work in basically the same way, they all use the same kind of workstation, and they all use the same visualization tools and software packages. All these factors make it easier to upgrade AEROnet to support the types of applications presented herein, than to enhance the performance of a general-purpose network. Nevertheless, experiences within this limited community should be useful in the more general context.

It is clear that today's networks do not offer the capabilities to support the future computational aerosciences applications that are presented in this paper. Although users with multiple T1 access to the NAS facility do not experience the frustrations expressed by those with only 56-kilobit/second access, their mode of operation is nevertheless restricted by current network limitations. Future plans for upgrade of AEROnet include the upgrade of some of the 56-kilobit/second lines to T1 and possibly upgrading the higher-speed links to T3. Eventual upgrade to gigabit rates is being considered, but is not feasible in the near future because of the expense.

We have shown that gigabit networks are necessary to enable interactive visualization and scientific collaboration. Although we have identified general networking requirements to support these applications, experimentation is needed to explore various protocols and implementation techniques to determine how to satisfy them. Several organizations within the Bay Area are planning to establish a Bay Area gigabit testbed as part of the HPCCP gigabit-network development. At NASA Ames we hope to use this testbed to experiment with implementing the types of applications that are described in Section 3 of this paper.

Acknowledgements. The author acknowledges the users at NASA LaRC, NASA LeRC, McDonnell Douglas, and Rockwell who so willingly provided insight into their experiences using the NAS facility. Also, the author acknowledges personnel within the Networking Group at NAS, who readily provided current information on the NAS networking system.

References

1. E. Authement, "AEROnet: A Network Built to Standards," */u/nas/news, The Numerical Aerodynamic Simulation Program Newsletter*, Vol. 6, No. 9, October 1991.
2. F. R. Bailey, "Research and Development in the Numerical Aerodynamic Simulation Program," in *Supercomputing*, Janusz S. Kowalik, editor, Springer-Verlag, 1990, pp. 55-69.
3. G. Chesson, *lensnoop*, program to monitor computer traffic.
4. D. Iannucci, "New Communication Utility to Increase Throughput," */u/nas/news, The Numerical Aerodynamic Simulation Program Newsletter*, Vol. 5, No. 6, June 1990.
5. J. Lekashman, "Computer Networking for Interactive Supercomputing," in *Supercomputing*, Janusz S. Kowalik, editor, Springer-Verlag, 1990, pp. 397-408.
6. E. Levin, C.K. Eaton and B. Young, "Scaling of Data Communications for an Advanced Supercomputer Network," *Data Communication Systems and Their Performance*, North Holland, 1988, pp. 107-122.
7. B. H. McCormick, T. A. DeFanti, and M. D. Brown (eds.), "Visualization in Scientific Computing," *Computer Graphics* 21 (6), Nov. 1987.
8. Office of Science and Technology Policy, Executive Office of the President, "Grand Challenges: High Performance Computing and Communications. FY 1992 U.S. Research and Development Program," 1991.
9. A. Vaziri, "Scientific Visualization in High-Speed Network Environments," *Computer Networks & ISDN Systems*, Vol. 22, No. 2 (1991), pp. 111-129.
10. R. G. Voigt, "Requirements for Multidisciplinary Design of Aerospace Vehicles on High Performance Computers," in *Supercomputing*, Janusz S. Kowalik, editor, Springer-Verlag, 1990, pp. 227-237.